# Linking CGE and Microsimulation Models for Distributional Policy Analysis

A Macro–Micro Framework with Twin Applications to Trade Liberalization and Carbon Pricing

Israel Osorio-Rodarte*

This draft: June 14, 2026

**Abstract**

Distributional policy analysis with computable general equilibrium (CGE) models turns on a methodological choice—how household heterogeneity enters the model—yet that choice is rarely interrogated by holding it fixed across very different policies. This paper does so. Using one transparent, reproducible macro–micro spine—a predominantly top-down coupling of the recursive-dynamic ENVISAGE CGE model to the World Bank's Global Income Distribution Dynamics (GIDD) microsimulation—it analyzes two policies that differ in almost every respect: trade liberalization and supply-chain reorganization (Chepeliev et al. 2025), and the poverty and distributional incidence of global carbon pricing (Osorio-Rodarte, Chepeliev, and Maliszewska 2025). The same spine recovers broad-based, mildly pro-poor gains from liberalization and a small, front-loaded, regionally concentrated cost from carbon pricing—a cost that reverses sign, from roughly five million more extreme poor to a net reduction, once the health co-benefits of cleaner air are valued. Neither incidence pattern is visible in representative-agent CGE results, and both are properties of the policy that the linkage makes legible. The paper situates the spine in the CGE–microsimulation literature, documents the global household-survey data foundations it rests on, and sketches where AI-assisted, provenance-tracked tooling could extend it.

**Keywords:** CGE; microsimulation; income distribution; poverty; trade; carbon pricing; ENVISAGE; GIDD.
**JEL:** C68, D31, F14, Q52, I32.

## 1 Introduction

Trade disruptions caused by recent shocks—the U.S.–China trade war, the COVID-19 pandemic, and geopolitical conflict—have renewed interest in reshoring and supply-chain diversification, even as the gains from openness remain substantial (Chepeliev et al. 2025). At the same time, the escalation of climate mitigation has put the distributional incidence of carbon pricing at the center of policy debate (Osorio-Rodarte, Chepeliev, and Maliszewska 2025). In both arenas, aggregate welfare measures are insufficient: policymakers need to know *who* gains and loses, and by how much. Capturing such heterogeneous impacts requires linking economywide computable general equilibrium (CGE) models—which deliver internally consistent, economywide responses to policy—with microsimulation models that map those responses onto the full distribution of households (Bourguignon and Bussolo 2013).

*World Bank. Corresponding author: `osorio.israel@gmail.com`. Prepared for the 29th Annual Conference on Global Economic Analysis (GTAP), Kyoto, June 17–19, 2026.

The difficulty is that the two model classes answer different questions at different resolutions. A CGE model tracks how a policy reshapes prices, factor returns, and output across sectors and regions, but typically collapses the household sector into a handful of representative agents, suppressing precisely the within-group heterogeneity that determines poverty and inequality outcomes. A microsimulation, conversely, preserves household detail but lacks a general-equilibrium account of where the shocks it imposes come from. Linking the two recovers both margins at once—provided the link is transparent about what is passed between the models, what is assumed, and where the data run out.

This paper has three aims. *First*, it documents the established ENVISAGE–GIDD macro–micro linkage as a reusable methodological framework, formalizing the vector of statistics passed from macro to micro and the wage-bill decomposition that keeps the two accounts consistent (§2). *Second*, it situates that framework within the broader literature on CGE–microsimulation integration (§3) and reviews the global household-survey data foundations on which it depends (§4). *Third*, it demonstrates the framework's reach through twin applications that share one methodological spine: trade liberalization and supply-chain reorganization (§5) and the poverty and distributional incidence of global carbon pricing (§6).

The contribution is deliberately methodological rather than topical. The two applications are chosen because they differ in almost every respect—one varies the openness of the world economy, the other prices an externality—yet are analyzed with an identical linkage. That the same spine recovers broad-based, mildly pro-poor gains in one case and a small, front-loaded, regionally concentrated cost in the other is the central demonstration: a representative-agent CGE, which holds the within-group distribution fixed by construction, cannot represent either incidence pattern, whereas the linkage recovers both. §7 draws out the limitations and sketches a research-infrastructure agenda—including AI-assisted, provenance-tracked tooling—for extending the approach; §8 concludes.

# 2 Linking Global CGE and Microsimulation Models

The framework links the ENVISAGE CGE model—a recursive-dynamic, multi-region model calibrated to the GTAP database (Mensbrugghe 2025)—to the Global Income Distribution Dynamics (GIDD) microsimulation (Bussolo, De Hoyos, and Medvedev 2010). ENVISAGE represents production through nested constant-elasticity-of-substitution functions and tracks factor accumulation over time; GIDD applies macro shocks to nationally representative household surveys to recover changes in the income distribution. The linkage is predominantly *top-down*—ENVISAGE transmits a compact vector of summary statistics to GIDD, which updates household incomes and consumption accordingly—but it is not strictly unidirectional. Some inputs originate on the micro side: the demographic and educational structure that anchors the baseline comes from the household data and population projections, and later applications of the framework replace the model's labor volumes and wages with earnings from the harmonized Gender Disaggregated Labor Database (World Bank 2024a), as in the African Continental Free Trade Area analysis (Echandi, Maliszewska, and Steenbergen 2022). The applications in this paper use the top-down configuration, which is appropriate where the distributional effects are modest relative to the aggregate shock (Bourguignon and Bussolo 2013).

## 2.1 The linkage vector

For each country, scenario, and year, ENVISAGE produces a compact set of transmission variables that summarize the macro shock as seen by the microdata—the channels through which prices and factor markets reach the household. These fall in three groups:

1. **Prices.** Laspeyres consumption-price indices by expenditure bundle —food, energy, transport, and the rest—and an aggregate index, mapped to household expenditure shares so that the cost-of-living shock varies with each household's consumption basket. Because those shares shift systematically across the distribution—food and energy together claim a far larger combined share of poorer households' budgets, driven mainly by the steep food gradient (Figure 1)—a bundle-specific price shock is mechanically regressive or progressive depending on which prices move.

2. **Factor markets.** Skill- and location-specific wage changes and the associated labor-supply shares, described in §2.2.

3. **Macro aggregates.** Real GDP per capita and per-capita consumption growth, which anchor the level of the distribution.

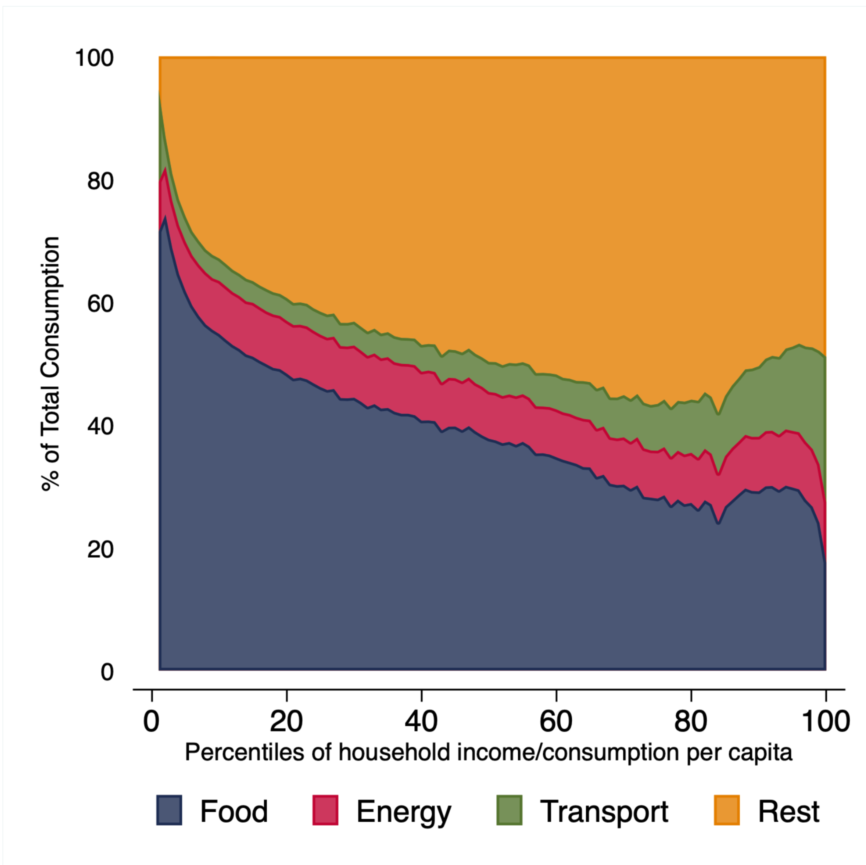


**Figure 1:** Household expenditure shares by bundle across the global distribution. Food and energy absorb a much larger share of poorer households' budgets, so the incidence of a bundle-specific price shock depends on which prices move. Source: (Osorio-Rodarte, Chepeliev, and Maliszewska 2025).

The microsimulation applies these shocks in a fixed sequence, stated here and formalized in Appendix A. (i) *Reweighting*: survey weights are rescaled by minimum cross-entropy to match projected age $\times$ education $\times$ gender population margins, building demographic and human-capital trends into the baseline before any policy shock (Bussolo, De Hoyos, and Medvedev 2010). (ii) *Sectoral reallocation*: in response to the CGE's labor-demand changes, workers move from contracting to expanding sectors *within* their skill–gender group, selected by an estimated multinomial-logit assignment; movers are assigned a wage in the receiving sector from a sector-specific Mincer equation, with the residual rescaled to preserve within-cell dispersion. (iii) *Wage-premium reconciliation*: wages are then scaled so that the *relative* skill–sector premia in the survey reproduce the proportional changes in those premia modeled by the CGE (Bourguignon and Bussolo 2013). (iv) *Income scaling*: a single scalar $\lambda$ rescales household income so that mean per-capita income growth matches the CGE, anchoring the *level* the premia reconciliation leaves free. (v) *Price adjustment*: real income is deflated by a household-specific food / non-food price index whose food share is estimated from the survey. Passing relative premia rather than nominal wages is what keeps the macro and micro accounts mutually consistent without forcing wage gaps to move in opposite directions across the two models.

## 2.2 The wage-premium reconciliation

Because labor income is the dominant income source for households near the poverty line, the fidelity of the linkage rests on how labor earnings are updated and reconciled to CGE factor returns. Each working-age earner in the survey is classified by *skill*—two levels, proxied by educational attainment (completed secondary or tertiary versus less)—and by *gender*, and is observed in one of the model's production *sectors*; the skill×gender classification, applied across sectors, is the unit at which wages are updated. This is the multi-sector extension of the original two-sector (agriculture / non-agriculture) GIDD design (Bussolo, De Hoyos, and Medvedev 2010), implemented in the model code and documented in full in Appendix A.

The update operates on relative wage *premia*, not nominal wages. Let $g_s = y_s/y_0 - 1$ be the CGE wage gap of skill–sector cell $s$ relative to the numeraire cell $y_0$ (unskilled agricultural labor) in the benchmark, and $\hat{g}_s$ the corresponding gap in the policy simulation; let $g'_s$ be the analogous premium computed from the survey. The counterfactual survey premium is

$$\hat{g}'_s \;=\; g'_s \,\frac{\hat{g}_s}{g_s},$$

so the survey inherits the CGE's *proportional* change in each premium while retaining its own baseline wage structure—which prevents the macro and micro wage gaps from drifting apart. A single scalar $\lambda$ then rescales all household incomes so that weighted mean per-capita income growth matches the CGE aggregate, fixing the level that the premia leave undetermined. Full expressions are given in Appendix A.

Two scope conditions follow from this design and bear directly on interpretation. *First*, the reconciliation operates on *labor* income only. Non-labor components—capital returns, land rents, remittances, and transfers—are not separately micro-updated; they move only with the common scalar $\lambda$. As the original derivation notes, this is a reasonable approximation for poor households, who hold little capital, but it introduces error higher in the distribution and means the framework is not suited to questions where asset income drives the result. *Second*, location (rural/urban) enters the model as a covariate in the sectoral-assignment step, not as a separate wage cell; wage premia are reconciled on the skill×gender×sector margin. Extending a consistent rural/urban earnings split across the full survey collection—and, with it, a spatially resolved wage update—remains a data-infrastructure priority, discussed in §7.

## 2.3 Assumptions and scope

Two simplifications bound the interpretation of results. First, labor markets clear each period and workers move across sectors at no cost, abstracting from adjustment frictions and unemployment that would amplify short-run distributional impacts. Second, the absence of micro-to-macro feedback means general-equilibrium effects of distributional change are not recirculated; this is defensible when those effects are second-order, as in the applications below, but would bind in settings with large transfers or behavioral responses.

# 3 Literature on Macro–Micro Linkages

With the linkage specified, it can be placed among the alternatives the literature has developed—both to justify the top-down choice and to mark what it gives up. The problem of recovering distribution from an economywide model is as old as applied general equilibrium itself. Early CGE models in the tradition of Dervis, Melo, and Robinson (1982) represented the household sector with a small number of representative agents, holding within-group inequality fixed by construction. Capturing distributional change therefore required coupling the macro model to disaggregated household data—a problem foreshadowed by early development-planning models (Adelman and Robinson 1978) and approached in several distinct ways over the following decades (Davies, St-Hilaire, and Whalley 1984; Bourguignon, Bussolo, and Pereira da Silva 2008).

## 3.1 A taxonomy of approaches

It is useful to organize the literature along two axes: *where* household heterogeneity lives (inside the CGE or in a separate microsimulation), and *whether* information flows in one direction or two (Table 1). At one pole, the *representative-household* (RH) approach keeps a few stylized agents inside the CGE; it is computationally light but silent on intra-group inequality. The *integrated multi-household* (IMH) approach embeds many—sometimes all—survey households directly in the CGE's social accounting matrix (Decaluwé et al. 1999; Cockburn 2006), giving full consistency at a high data and computational cost. *Top-down* (or "soft-linked") approaches pass prices and factor returns from the CGE to a standalone microsimulation without feedback—the lineage of sequential microsimulation models applied to developing economies (Cogneau and Robilliard 2000; Bourguignon, Robilliard, and Robinson 2003; Colombo 2010); while *top-down / bottom-up* approaches iterate the two to convergence, restoring feedback at additional cost (Savard 2003; Robilliard, Bourguignon, and Robinson 2008). Cross-cutting this, the microsimulation itself may be *arithmetic* (accounting-based reweighting and income mapping, as in GIDD) or *behavioral* (estimating labor-supply or occupational-choice responses)—a distinction that governs how much household adjustment the model endogenizes.

**Table 1:** Approaches to distribution in CGE modeling.

| Approach | Households | Feedback | Data / compute | Representative work |
|---|---|---|---|---|
| Representative household | In CGE (few) | — | Low | Dervis, Melo, and Robinson ( |
| Integrated multi-household | In CGE (many) | Full (by construction) | High | Decaluwé et al. (1999); Cockb |
| Top-down (soft link) | Separate microsim | None | Moderate | Bourguignon, Robilliard, and |
| Top-down / bottom-up | Separate microsim | Iterated | High | Savard (2003) |

## 3.2 Where this framework sits

ENVISAGE–GIDD is a top-down, arithmetic microsimulation: it inherits the computational tractability and global scale that a behavioral or fully integrated model would sacrifice, at the cost of suppressing micro-to-macro feedback and household behavioral response. Within the GTAP tradition, the dominant route to distribution has been the *stratified-household* approach, which partitions each region's representative household into earnings strata inside the CGE (Hertel et al. 2004; Ivanic and Martin 2008; Anderson, Cockburn, and Martin 2010); the present linkage instead passes the macro shock to the full survey distribution, trading that approach's in-model consistency for the household-level granularity a survey retains. Colombo (2010) shows that the choice among these methods materially affects results, with differences driven by data consistency and the presence of feedback; Peichl (2008) documents that linked models exploit complementary macro and micro strengths but that the linking step is itself non-trivial and consequential. The applications below sit squarely in the regime where the top-down simplification is defensible—global analyses in which the distributional effects are modest relative to the aggregate shock, so that omitted feedback is second-order (Hertel and Winters 2006). Country-scale studies, by contrast, often find large and heterogeneous incidence that a single-economy integrated model is better placed to capture (Boccanfuso and Savard 2006). Throughout, distributional outcomes are read against standard global poverty lines and measurement conventions (Chen and Ravallion 2010).

# 4 Data Foundations for Global Microsimulation

A global microsimulation is only as credible as the household microdata beneath it. GIDD draws on a harmonized collection of roughly 130 nationally representative household surveys; across the collection these contain on the order of 10.5 million sampled *individuals* (person records) and cover about 83 percent of global GDP. The specific vintage used for the applications here comprises 11 621 774 individual records, whose expanded (sample-weighted) population covers about 88.87 percent of the world's people in 2019 (Table 2, where "Survey obs." counts person records). Coverage is high in most regions but structurally thinner in the Middle East and North Africa, where harmonized survey access is more limited.

**Table 2:** GIDD survey coverage by region. Population in millions; coverage as a share of regional population. Source: (Osorio-Rodarte, Chepeliev, and Maliszewska 2025).

| Region | Survey obs. | Pop. 2019 (m) | Coverage 2019 (%) |
|---|---|---|---|
| East Asia & Pacific | 1 010 643 | 1978.7 | 86.3 |
| Europe & Central Asia | 891 620 | 801.6 | 92.1 |
| Latin America & Caribbean | 2 515 162 | 592.2 | 92.3 |
| Middle East & North Africa | 690 785 | 233.0 | 52.8 |
| North America | 3 679 435 | 358.2 | 100.0 |
| South Asia | 954 037 | 1744.4 | 98.0 |
| Sub-Saharan Africa | 1 880 092 | 935.1 | 85.9 |
| World | 11 621 774 | 6643.2 | 88.9 |

Several complementary collections can broaden coverage and enrich the variable set. The Global Labor Database harmonizes labor-force and household surveys with detailed employment and earnings detail (World Bank 2024c); the Global Micro Database adds harmonized surveys with comparatively strong low-income coverage, particularly in Sub-Saharan Africa (World Bank 2024b). The Luxembourg Income and Wealth Studies provide harmonized income and—distinctively—wealth microdata for dozens of mostly high-income countries (LIS Cross-National Data Center 2024); SEDLAC covers Latin America and the Caribbean (CEDLAS and World Bank 2024); and EU-SILC offers harmonized European data (Eurostat 2024). IPUMS International contributes census microdata for over 100 countries (Ruggles et al. 2015), valuable for demographic and locational detail though thin on income and consumption aggregates. Sectoral and modeling systems such as IFPRI's RIAPA extend the toolkit toward agriculture and food systems (International Food Policy Research Institute 2025).

Two data constraints bear directly on the method. First, the skill and location detail required for the wage-bill decomposition of §2.2 is unevenly available: educational attainment supports a robust skilled/unskilled split almost everywhere, but a consistent rural/urban earner classification is present in only a subset of surveys. Second, integrating heterogeneous collections—through statistical matching or synthetic-population techniques—can fill geographic and variable gaps, but risks distorting the joint distribution of demographics, skills, income, and consumption that makes microsimulation informative. Observed microdata therefore remain the preferred foundation, with synthetic and matched data used to complement rather than substitute for them.

These data underpin a baseline against which every policy scenario is measured. Global extreme poverty fell from 45.2 percent in 1990 to single digits by the late 2010s; the recursive-dynamic baseline carries it to about 6.2 percent by 2050—above the 3 percent extreme-poverty target of Sustainable Development Goal 1 (Figure 2). How that baseline is projected, however, is not neutral. Figure 3 contrasts three treatments that differ in *how much* of the microsimulation machinery is switched on: (i) *reweighting only*, which rescales survey weights to projected age–education–gender margins but leaves every household's income growing at the common macro rate; (ii) *distribution-neutral growth*, which scales all incomes by a single factor to match macro per-capita growth with no change in relative positions; and (iii) the *full GIDD*, which adds the sector– and skill-specific wage and price effects of §2 on top of reweighting. (Reweighting is thus a component of full GIDD, not a rival to it; the first treatment isolates that component.) The three produce visibly different 2050 poverty paths—previewing the paper's central point that the treatment of household heterogeneity, not just the macro projection, drives the distributional result.

# 5 Application I: Trade Liberalization and Supply-Chain Reorganization

With the linkage specified (§2), positioned in the literature (§3), and its data foundations laid out (§4), it can be put to work on two contrasting policies. The first varies the openness of the world economy.

The first application uses the ENVISAGE–GIDD framework to evaluate trade-policy scenarios relative to a post-COVID baseline (Chepeliev et al. 2025). Three scenarios are examined: (i) *Reshoring Leading Economies*, in which high-income countries subsidize domestic production, impose import surcharges, and lower trade elasticities; (ii) *Reshoring All*, in which all countries adopt similar protectionist policies; and (iii) a *GVC-friendly liberalization*, in which developing countries eliminate tariffs on intermediate inputs, relax product standards, and adopt trade-facilitation measures. A supply-chain shock to Thailand's electronics sector is simulated under each to gauge how exposure to a localized disruption varies with trade regime.

Reshoring is contractionary and regressive. When leading economies reshore, world real income falls by about 1.5 percent, with the largest losses in trade-dependent developing economies—roughly 8.5 percent in Vietnam and 4.4 percent in Thailand. When all countries reshore, the global loss widens to about 2.2 percent and Sub-Saharan Africa contracts by some 5.5 percent. The poverty consequences track these losses: reshoring by all countries pushes about 51.8 million more people into extreme poverty (at the $1.90 PPP line used by the source study) by 2030, with Sub-Saharan Africa absorbing roughly 80 percent of the new poor. The incidence is regressive within regions as

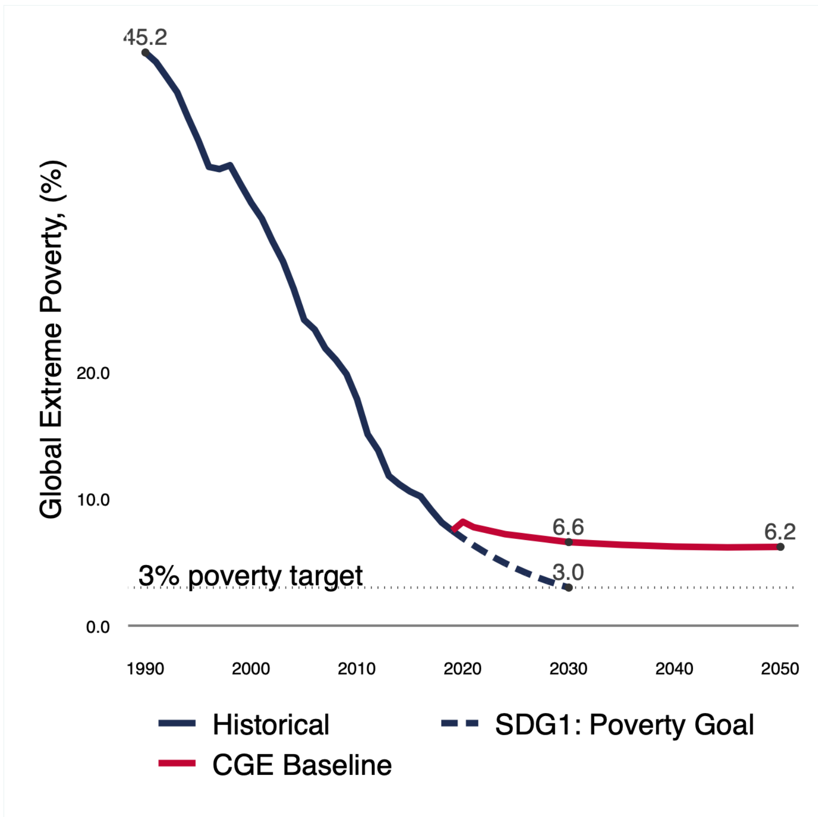


**Figure 2:** Global extreme poverty ($2.15 PPP), 1990–2050: the historical series, the path implied by the Sustainable Development Goal 1 target of 3 percent, and the recursive-dynamic CGE baseline. Source: (Osorio-Rodarte, Chepeliev, and Maliszewska 2025).

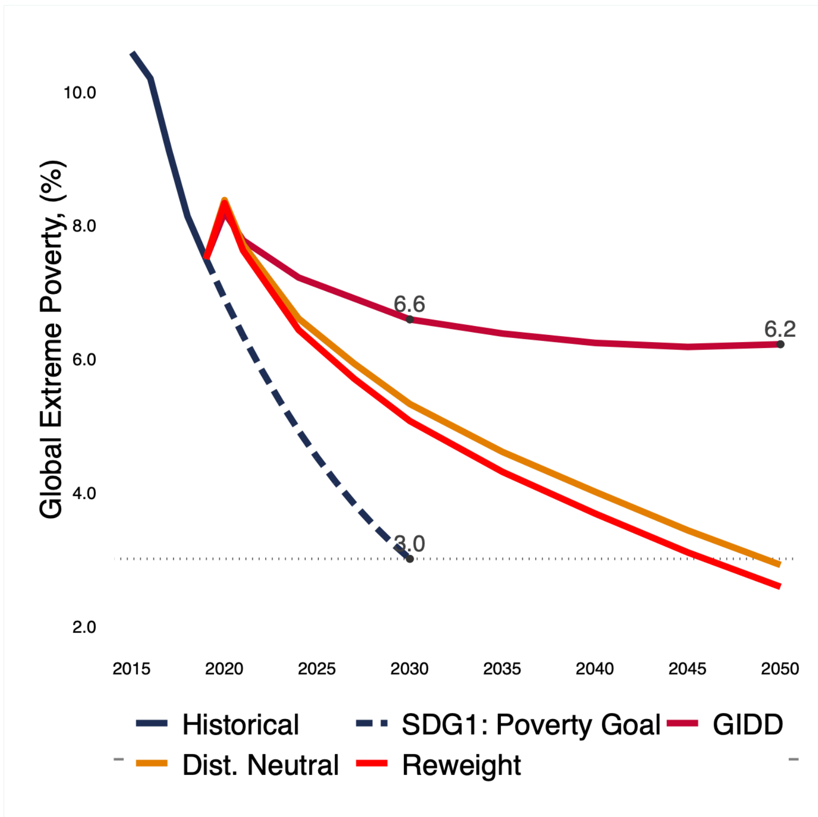


**Figure 3:** Baseline poverty projections under three treatments that switch on successively more of the microsimulation: reweighting only (demographic margins), distribution-neutral growth (uniform income scaling), and full GIDD (reweighting plus sector/skill wage and price effects), shown against the historical series and the SDG 1 target. The spread isolates how much the projected distribution depends on the microsimulation treatment, not the macro projection alone. Source: (Osorio-Rodarte, Chepeliev, and Maliszewska 2025).

well—in Europe and Central Asia, the bottom 40 percent lose about 8.8 percent of income under Reshoring All, against 1.7 percent for the top 60 percent.

Liberalization reverses the pattern. The GVC-friendly scenario lowers trade costs and raises real incomes across developing regions—gains on the order of 10.7 percent in Thailand and 6.8 percent in Vietnam—and reduces global extreme poverty (same $1.90 line) by about 21.5 million by 2030. Gains are pro-poor: in Europe and Central Asia the bottom 40 percent see welfare rise about 3.9 percent relative to baseline. Real incomes rise across the distribution, with larger proportional gains lower down.

The applications also illustrate how exposure to a localized disruption varies with trade regime. At the aggregate level, Thailand's real-income loss from the simulated electronics supply-chain shock is about 26 percent smaller under GVC-friendly liberalization than under Reshoring All: diversified, open supply chains absorb a localized disruption that concentrated, reshored ones amplify. That aggregate comparison is itself a CGE outcome—visible without any microsimulation. What the GIDD linkage adds is the *distributional* face of the same shock: how the loss is shared across the income distribution, and hence how much of it lands on the poor under each trade regime. This result is illustrative rather than general: it rests on a single shock (one sector, one country, two scenarios), and the frictionless labor reallocation assumed in §2 mechanically eases the open economy's adjustment, so the magnitude should be treated as suggestive pending a systematic sweep of shocks.

# 6 Application II: Carbon Pricing and Distribution

The second application turns the same ENVISAGE–GIDD spine on the distributional incidence of global carbon pricing through 2050 (Osorio-Rodarte, Chepeliev, and Maliszewska 2025). Where the trade application (§5) varies the openness of the world economy, here the policy lever is a carbon price path, and the question is who bears the transition's cost as energy and food prices rise. The linkage is identical in form—percentage changes in consumption prices by bundle, in employment, and in skill-specific wage premia are passed from the macro model to the household microdata, with wage premia scaled so that the microsimulation reproduces the relative-wage movements modeled by the CGE (Bourguignon and Bussolo 2013)—only the shock differs.

## 6.1 Scenarios

Seven policy scenarios are evaluated against a no-new-policy baseline, with carbon revenues recycled domestically as lump-sum transfers (Table 3). Five are carbon-price scenarios spanning a gradient of ambition, from meeting current Nationally Determined Contributions (NDC) to a 1.5 °C–consistent path that prices carbon already in the 2020s (<2 °C-Early). The remaining two of the seven (NDC-CO, 2 °C-CO) repeat the NDC and 2 °C cases but additionally monetize the health co-benefits of reduced air pollution—avoided premature mortality valued at country-specific, income-adjusted statistical lives—returned to households as a uniform per-capita income gain.

## 6.2 Poverty

Carbon pricing raises extreme poverty (the $2.15 PPP line) relative to baseline, but the magnitude is modest. By 2050 the NDC pathway leaves about 5.1 million more people in extreme poverty; adding the EU border adjustment brings this to 5.7 million. More ambitious pricing raises the figure to 6.0 million under 2 °C and 7.3 million under <2 °C-Early. In headcount terms these are increases of only 0.06–0.09 percentage points in the global poverty rate.

**Table 3:** Carbon-pricing scenarios.

| Scenario | Definition |
|---|---|
| Baseline | Current policies only; no new globally coordinated carbon price. |
| NDC | 2030 NDC targets met; moderate carbon pricing thereafter to hold the trajectory. |
| NDC-CBAM | NDC plus an EU carbon border adjustment on carbon-intensive imports. |
| 2 °C | Harmonized global carbon price post-2030 on a 2 °C–consistent budget. |
| <2 °C | Steeper post-2030 pricing on a tighter budget (~1.8 °C by 2100). |
| <2 °C-Early | Coordinated pricing begun in the 2020s; 1.5 °C–consistent. |
| NDC-CO, 2 °C-CO | As NDC / 2 °C, with monetized health co-benefits added to incomes by 2050. |

The transition cost is front-loaded. Because the price shock is sharpest while economies adjust to higher energy costs, poverty impacts peak in the 2030s and recede thereafter. Under 2 °C the additional extreme poor number 4.1 million in 2030 but 11.3 million by 2040; the <2 °C-Early path, which front-loads mitigation, spikes to 23.1 million in 2030 before falling back to 13.7 million by 2040. The near-term, not the 2050 endpoint, is where the burden is heaviest.

These increases are geographically concentrated. India and Sub-Saharan Africa together account for most of the global rise: under NDC, India adds roughly 3.0 million extreme poor by 2050 and Sub-Saharan Africa roughly 2.1 million. Both regions combine large poor populations near the poverty line with high exposure to energy-price increases and limited fiscal space for compensation.

## 6.3 The co-benefit sign reversal

Accounting for health co-benefits reverses the sign of the poverty effect. When the monetized value of cleaner air is returned to households, the NDC pathway *reduces* extreme poverty by about 2.7 million relative to baseline by 2050, and the 2 °C pathway by about 5.7 million—swings of roughly 7.8 and 11.8 million people, respectively, relative to the same scenarios without co-benefits (Figure 4, which plots poverty *increases* as downward bars and *reductions* as upward bars). Because the co-benefit is modeled as a uniform per-capita income gain, the *sign* of the reversal at a given horizon follows mechanically once the transfer exceeds the underlying poverty increase; the substantive result is the *magnitude* of the monetized health dividend, which by 2050 is large enough to outweigh the carbon-price drag (co-benefit scenarios are evaluated at the 2050 horizon). The distributional case for mitigation thus hinges on whether its health dividend is counted: priced on energy costs alone it is mildly regressive in the aggregate, but valued inclusive of avoided mortality it is poverty-reducing.

## 6.4 Inequality

Effects on global inequality are small and reflect offsetting movements. Carbon pricing modestly raises *between*-country inequality—poorer countries bear proportionally larger losses—while remaining close to neutral *within* countries: the global Gini stands at 66.06 under 2 °C in 2045, essentially unchanged from baseline. Within most regions the growth-incidence pattern is flat or mildly progressive: in the European Union, United States, China, and Indonesia, higher-income quintiles absorb the larger proportional income declines, leaving the bottom of the distribution near-neutral. India is the notable exception: in the 2 °C and <2 °C-Early pathways its lower quintiles bear proportionally larger losses—a regressive within-country pattern that matters because India is also one of the largest contributors to the global poverty increase (§6.2). The growth-incidence curves (Figure 5) make both the general pattern and the Indian exception concrete.

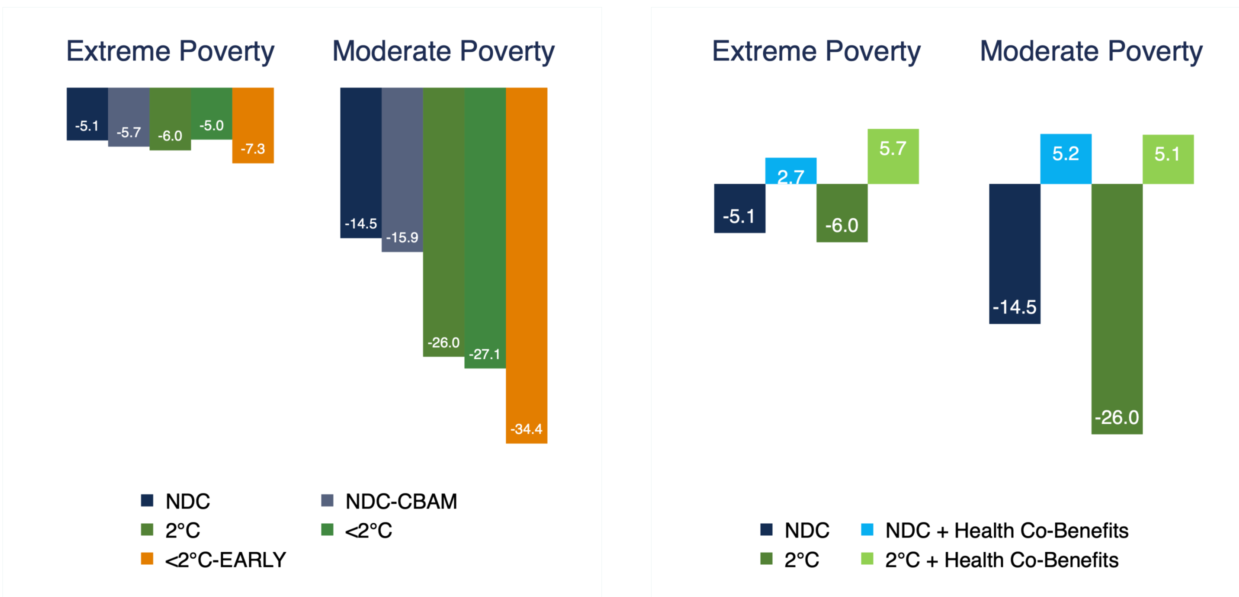


(a) Carbon pricing only

(b) With monetized health co-benefits

**Figure 4:** Change in the global poor by 2050, by scenario, without (a) and with (b) monetized health co-benefits. Each panel reports two poverty lines: extreme poverty ($2.15 PPP, left) and moderate poverty ($3.65 PPP, right); the text discusses the extreme-poverty results. Bars are plotted in a welfare convention: *downward (negative) bars denote an increase in poverty*, upward (positive) bars a reduction. Thus in panel (a) every carbon-pricing scenario lies below zero (more poverty) at both lines, while in panel (b) the co-benefit scenarios cross above zero (fewer poor than baseline)—the sign reversal discussed in the text. Magnitudes are millions of people. Source: (Osorio-Rodarte, Chepeliev, and Maliszewska 2025).

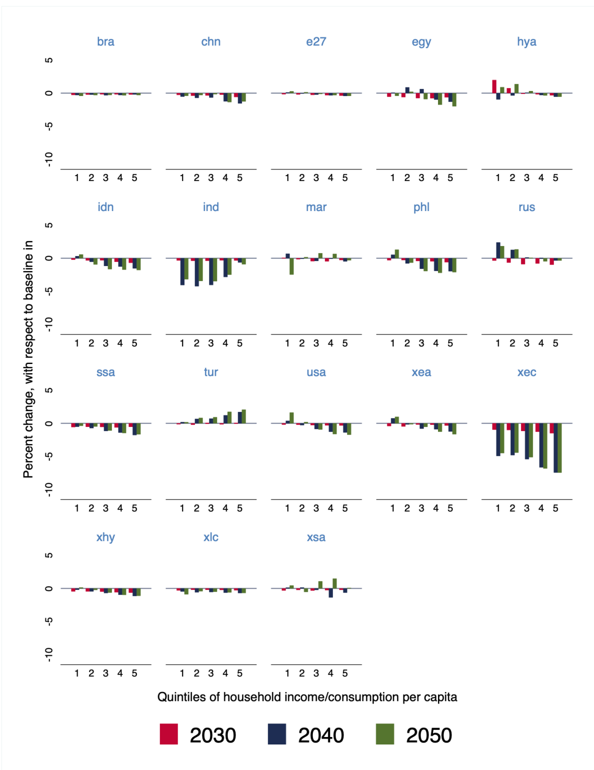


(a) 2 °C

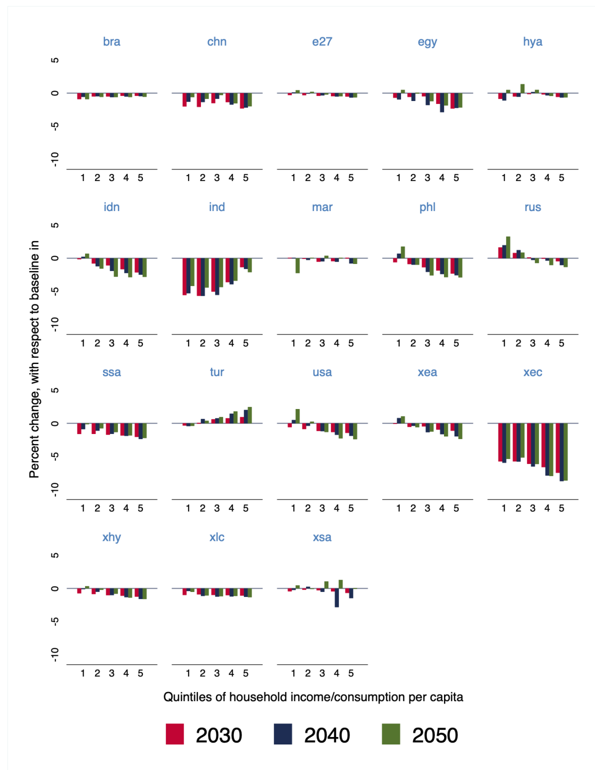


(b) <2 °C-Early

**Figure 5:** Growth-incidence bars by region and global income quintile, relative to baseline. In most regions the pattern is flat-to-progressive, indicating roughly neutral within-country incidence, though some regions show larger uniform losses. Source: (Osorio-Rodarte, Chepeliev, and Maliszewska 2025).

### 6.5 Reading across the two applications

Put the two applications side by side. The linkage delivers broad-based, mildly pro-poor gains under trade liberalization and a small, front-loaded, regionally concentrated cost under carbon pricing—and in the second case the cost changes sign once health co-benefits are counted. Two incidence patterns this different share one spine; a model that fixes the within-group distribution sees neither, because both live entirely in the household detail the linkage restores. (The two applications report against different poverty lines—\$1.90 for trade, \$2.15 for carbon—inherited from their source studies; the contrast is in the shape of incidence, not a level comparison across the two.)

## 7 Discussion and Research Directions

The twin applications illustrate both the reach and the limits of a soft, top-down macro–micro linkage. This section is candid about the second before turning to where the approach can go next.

### 7.1 Limitations

Three simplifications bound the results. *No micro-to-macro feedback*: the one-way link means distributional changes do not recirculate into prices or factor markets. This is defensible when distributional effects are second-order, as in both applications, but would bind under large transfers or sharp behavioral responses. *Stylized household behavior*: the arithmetic microsimulation reallocates labor by similarity and imputes wages with Mincer equations, but does not estimate labor-supply or occupational-choice elasticities, so behavioral adjustment is muted and short-run frictions—adjustment costs, unemployment—are absent. *Simplified incidence of some channels*: the carbon co-benefits, for instance, are returned to households as a uniform per-capita gain rather than modeled with household-specific exposure to air pollution—a simplification that likely understates heterogeneity in who benefits. Each of these is a deliberate trade for global scale and transparency, and each marks a direction for refinement.

### 7.2 Extending the framework

Several extensions follow directly. Dynamic microsimulation could track life-cycle decisions, human-capital formation, and demographic change, opening long-run distributional questions the recursive baseline only approximates. Iterative top-down/bottom-up coupling would restore the feedback the current link omits. Richer data integration—across the Luxembourg studies, IPUMS, and the labor and micro databases of §4—could deepen coverage and, by harmonizing a consistent skill×location earner classification, close the gap between the implemented and target wage-bill decompositions of §2.2. Multi-model linkage is a further frontier: the ENVISAGE–GIDD chain has already been augmented with monetized health co-benefits derived from air-pollution modeling (Osorio-Rodarte, Chepeliev, and Maliszewska 2025), and analogous links to labor-market, nutrition, or climate-impact models would extend the reach of the spine.

### 7.3 A research-infrastructure agenda

A distinct frontier is methodological infrastructure rather than model structure. Macro–micro analysis of the kind set out here is bottlenecked less by economic theory than by the labor of wiring models together reproducibly: extracting the linkage vector, harmonizing survey variables, propagating shocks, and—above all—keeping every reported number traceable to the data cell that produced it. Tooling that treats this plumbing as a first-class object is under development.

The design organizes the workflow around provenance: each quantitative claim in a deliverable carries a reference back to a specific (variable, region, scenario, year, value) record, and an automated coherence check verifies that aggregate (CGE) and distributional (micro) narratives remain mutually consistent as inputs change. Reproducibility is enforced by pinning model and tool versions and by validating generated output against hand-checked baselines.

There is scope here for AI-assisted methods—in particular, agent-based orchestration that decomposes the macro→micro storyline into checkable steps and flags inconsistencies for human review. This is framed deliberately as research infrastructure, with a human economist in the loop: the aim is to make distributional analysis faster to assemble and easier to audit, not to automate judgment. The agenda is at an early, pre-deployment stage—a direction under active development rather than a result—and the present paper's findings stand independently of it.

# 8 Conclusion

Linking CGE and microsimulation models lets researchers move beyond aggregate welfare to recover the distributional effects of policy reform. Applied to trade, the ENVISAGE–GIDD framework shows that reshoring is contractionary and regressive—pushing tens of millions into poverty, four-fifths of them in Sub-Saharan Africa—while GVC-friendly liberalization is pro-poor and, in a single supply-chain shock simulation, less exposed to a localized disruption. Applied to carbon pricing, the same framework shows a small, front-loaded increase in extreme poverty that is concentrated in South Asia and Sub-Saharan Africa and that reverses sign once the health co-benefits of cleaner air are counted. The point is not either result on its own but that a single macro–micro spine, held fixed across two very different policies, recovers incidence that representative-agent analysis—which assumes away within-group distribution—cannot represent. Continued progress in dynamic microsimulation, integrated modeling, survey harmonization, and reproducible research infrastructure will determine how far that spine can be pushed.

# 9 References

Adelman, Irma, and Sherman Robinson. 1978. *Income Distribution Policy in South Korea.* Stanford, CA: Stanford University Press.

Anderson, Kym, John Cockburn, and Will Martin, eds. 2010. *Agricultural Price Distortions, Inequality, and Poverty.* GTAP-linked global trade-distortion poverty analysis. Washington, DC: World Bank.

Boccanfuso, Dorothée, and Luc Savard. 2006. *Impacts Analysis of Cotton Subsidies on Poverty: A CGE Macro-Accounting Approach.* GREDI Working Paper 06-04. Working-paper version; later published in Journal of African Economies 16(4):629–659, 2007. GREDI, Université de Sherbrooke.

Bourguignon, François, and Maurizio Bussolo. 2013. "Income Distribution in Computable General Equilibrium Modeling." In *Handbook of Computable General Equilibrium Modeling,* edited by Peter B. Dixon and Dale W. Jorgenson, 1:1383–1437. Amsterdam: Elsevier.

Bourguignon, François, Maurizio Bussolo, and Luiz A. Pereira da Silva. 2008. *The Impact of Macroeconomic Policies on Poverty and Income Distribution: Macro-Micro Evaluation Techniques and Tools.* Washington, DC: World Bank / Palgrave Macmillan.

Bourguignon, François, Anne-Sophie Robilliard, and Sherman Robinson. 2003. "Representative versus Real Households in the Macroeconomic Modeling of Inequality." In *DELTA Working Paper 2003-05.* Layered/sequential macro-micro approach.

Bussolo, Maurizio, Rafael E. De Hoyos, and Denis Medvedev. 2010. "Economic Growth and Income Distribution: Linking Macro-Economic Models with Household Survey Data at the Global Level." VERIFY: confirm volume/issue/pages for the GIDD framework cite, *International Journal of Microsimulation* 3 (1): 92–103.

CEDLAS and World Bank. 2024. "Socio-Economic Database for Latin America and the Caribbean (SEDLAC)." CEDLAS, Universidad Nacional de La Plata and World Bank. Accessed January 10, 2026. https://www.cedlas.econo.unlp.edu.ar/wp/en/estadisticas/sedlac/.

Chen, Shaohua, and Martin Ravallion. 2010. "The Developing World Is Poorer Than We Thought, but No Less Successful in the Fight against Poverty." Global poverty-line methodology underpinning poverty measurement, *Quarterly Journal of Economics* 125 (4): 1577–1625.

Chepeliev, Maksym, Maryla Maliszewska, Israel Osorio-Rodarte, Maria Filipa Seara e Pereira, and Dominique van der Mensbrugghe. 2025. "Lowering Trade Barriers Improves Income Distribution and Economic Resiliency." *Journal of Policy Modeling* 47:30–48.

Cockburn, John. 2006. "Trade Liberalisation and Poverty in Nepal: A Computable General Equilibrium Micro Simulation Analysis." In *Globalisation and Poverty: Channels and Policy Responses,* edited by John Cockburn and Paolo Giordano. Integrated CGE-microsimulation applied to trade; widely cited top-down/integrated reference. London: Routledge.

Cogneau, Denis, and Anne-Sophie Robilliard. 2000. "Income Distribution, Poverty and Growth in a Developing Country: A Micro-Simulation Model Applied to Madagascar." Early integrated macro-micro distributional model, *DIAL Working Paper.*

Colombo, Giulia. 2010. "Linking CGE and Microsimulation Models: A Comparison of Different Approaches." Citekey retains 2018 label; actual publication year is 2010 (IJM 3(1)), *International Journal of Microsimulation* 3 (1): 72–91.

Davies, James B., France St-Hilaire, and John Whalley. 1984. "Some Calculations of Lifetime Tax Incidence." In *The American Economic Review,* 74:633–649. 4. Microsimulation–CGE incidence; cited as a seminal integrated-approach reference.

Decaluwé, Bernard, André Patry, Luc Savard, and Erik Thorbecke. 1999. "Poverty Analysis within a General Equilibrium Framework." In *CRÉFA Working Paper 9909, Université Laval.*

Dervis, Kemal, Jaime de Melo, and Sherman Robinson. 1982. *General Equilibrium Models for Development Policy.* Cambridge: Cambridge University Press.

Echandi, Roberto, Maryla Maliszewska, and Victor Steenbergen. 2022. *Making the Most of the African Continental Free Trade Area: Leveraging Trade and Foreign Direct Investment to Boost Growth and Reduce Poverty.* Washington, DC: World Bank.

Eurostat. 2024. "European Union Statistics on Income and Living Conditions (EU-SILC)." Eurostat. Accessed January 10, 2026. https://ec.europa.eu/eurostat/web/microdata/european-union-statistics-on-income-and-living-conditions.

Hertel, Thomas W., Maros Ivanic, Paul V. Preckel, and John A. L. Cranfield. 2004. “The Earnings Effects of Multilateral Trade Liberalization: Implications for Poverty.” GTAP stratified-household poverty analysis, *The World Bank Economic Review* 18 (2): 205–236.

Hertel, Thomas W., and L. Alan Winters, eds. 2006. *Poverty and the WTO: Impacts of the Doha Development Agenda.* GTAP-based trade-and-poverty distributional analysis. Washington, DC: World Bank / Palgrave Macmillan.

International Food Policy Research Institute. 2025. “RIAPA Data and Modeling System.” IFPRI. Accessed January 10, 2026.

Ivanic, Maros, and Will Martin. 2008. “Implications of Higher Global Food Prices for Poverty in Low-Income Countries.” GTAP-based food-price poverty incidence across stratified households, *Agricultural Economics* 39:405–416.

LIS Cross-National Data Center. 2024. “Luxembourg Income Study (LIS) and Luxembourg Wealth Study (LWS) Databases.” LIS. Accessed January 10, 2026. https://www.lisdatacenter.org/.

Maliszewska, Maryla, Israel Osorio-Rodarte, and Rakesh Gupta. 2020. *Ex-Ante Evaluation of Sub-National Labor Market Impacts of Trade Reforms.* Policy Research Working Paper 9478. GIDD microsimulation methodology; Annex 2 gives the linkage equations. Washington, DC: World Bank.

Mensbrugghe, Dominique van der. 2025. *The Environmental Impact and Sustainability Applied General Equilibrium (ENVISAGE) Model, Version 11.* Technical Reference Document. Center for Global Trade Analysis, Purdue University, September.

Osorio-Rodarte, Israel, Maksym Chepeliev, and Maryla Maliszewska. 2025. “The Poverty and Distributional Impacts of Carbon Pricing: A Global View.” World Bank, mimeo.

Peichl, Andreas. 2008. *The Benefits of Linking CGE and Microsimulation Models: Evidence from a Flat Tax Analysis.* IZA Discussion Paper 3715. Bonn: Institute for the Study of Labor (IZA).

Robilliard, Anne-Sophie, François Bourguignon, and Sherman Robinson. 2008. “Examining the Social Impact of the Indonesian Financial Crisis Using a Macro-Micro Model.” Originally circulated 2001; foundational microsimulation-sequential CGE application, *Poverty, Inequality and Policy in Latin America.*

Ruggles, Steven, Miriam King, Sarah Flood, Katie Genadek, Diana Magnuson, Jose Pacas, and Matthew Sobek. 2015. “IPUMS International: Harmonized International Census Microdata.” *Historical Methods* 48 (4): 195–206.

Savard, Luc. 2003. *Poverty and Inequality Analysis within a CGE Framework: A Comparative Analysis of the Representative Agent and Micro-Simulation Approaches.* PEP Working Paper. Working-paper version; later published in Development Policy Review 23(3):313–331, 2005. Poverty and Economic Policy (PEP) Research Network.

World Bank. 2024a. “Gender Disaggregated Labor Database (GDLD).” World Bank. Accessed June 14, 2026. https://datatopics.worldbank.org/gdld/.

———. 2024b. “Global Micro Database.” Poverty and Equity Global Practice. Accessed January 10, 2026.

———. 2024c. “Introduction to the Global Labor Database.” Global Labor Database. Accessed January 10, 2026. https://worldbank.github.io/gld/.

# A The GIDD algorithm

This appendix sets out the GIDD microsimulation in full, so that an expert reader can reconstruct the exact ENVISAGE-to-GIDD transformation. The core equations follow the original derivation of Bussolo, De Hoyos, and Medvedev (2010) and its statement in Maliszewska, Osorio-Rodarte, and Gupta (2020); the multi-sector skill×gender extension in §A.6 documents the version implemented in the model code released with this paper (§A.7).

## A.1 Income generation and the two sources of distributional change

Let the income (or consumption) $y_{i,t}$ of individual $i$ in year $t$ be a function of endowments $X_{i,t}$, the market rewards to those endowments $\beta_t$, an occupational/participation status $L_{i,t} \,|\, \lambda_t$, and an unobservable $\varepsilon_{i,t}$:

$$y_{i,t} \;=\; f(X_{i,t},\, \beta_t,\, (L_{i,t} \mid \lambda_t),\, \varepsilon_{i,t})\,. \tag{1}$$

The income distribution $D_t = \{y_{1,t}, \ldots, y_{N,t}\}$ then changes through two channels: changes in the *rewards* $\beta$ (and occupational parameters $\lambda$), and changes in the *distribution of endowments* $X$ (the age, education, and household-composition structure of the population). GIDD micro-simulates both.

For exposition with two education levels and two sectors, baseline income is

$$y_{i,t} = \alpha_t + \beta_{1,t}\Big(E^{S}_{i,t}S^{A}_{i,t}\Big) + \beta_{2,t}\Big(E^{S}_{i,t}S^{NA}_{i,t}\Big) + \beta_{3,t}\Big(E^{US}_{i,t}S^{NA}_{i,t}\Big) + \sum_{k=1}^{K}\gamma_{k,t}F_{k,t} + \varepsilon_{i,t}, \tag{2}$$

where $E^S, E^{US}$ are skilled/unskilled dummies, $S^A, S^{NA}$ are agricultural / non-agricultural sector dummies, $F_k$ is the share of household members in age cohort $k$, the $\beta$'s are sector-specific returns to skill, and the $\gamma$'s price household composition. The counterfactual $\hat{y}_{i,t+k}$ replaces the endowments $(E, S, F)$ and the returns $(\beta, \gamma)$ with their micro-simulated values; these are the *linkage aggregate variables* that connect GIDD to the CGE.

## A.2 Step 1 — demographic reweighting

Each country's population is partitioned into 16 five-year age groups, two genders, and three education levels (none/primary, secondary, tertiary). Age and gender totals follow the UN medium-variant projections; educational attainment evolves by a pipeline assumption (younger, more-educated cohorts replace older ones, with enrolment held at its base-year value). Survey weights are then recalibrated to these projected margins by minimum cross-entropy, i.e. the new weights deviate as little as possible from the originals subject to matching the targets. This is the "reweighting only" baseline treatment of Figure 3.

## A.3 Step 2 — wage-premium reconciliation

Let $y_{s,e}$ be average CGE earnings in sector $s$ and skill group $e$, with unskilled agricultural labor $y_{1,1}$ as numeraire. Define benchmark and counterfactual *wage gaps*

$$g_{s,e} = \frac{y_{s,e}}{y_{1,1}} - 1, \qquad \hat{g}_{s,e} = \frac{\hat{y}_{s,e}}{\hat{y}_{1,1}} - 1, \tag{3}$$

and the analogous *micro* premia from the household data, $g'_{s,e} = y'_{s,e}/y'_{1,1} - 1$. The counterfactual micro premium is obtained by passing the CGE's proportional change in the gap through to the survey,

$$\hat{g}'_{s,e} \;=\; g'_{s,e}\,\frac{\hat{g}_{s,e}}{g_{s,e}}. \tag{4}$$

Equation (4) is the heart of the linkage: even when the CGE and survey disagree on the *level* of a wage gap, the *percentage change* in each gap is made consistent across the two models, ruling out gaps that move in opposite directions. Updating operates on labor income only (eq. (2)); a subsequent scaling restores aggregate consistency.

## A.4 Step 3 — income scaling to the macro aggregate

Let $Y' = \sum_h \upsilon_h y_h$ and $\hat{Y}' = \sum_h \omega_h \widehat{y}_h$ be the weighted mean per-capita incomes in the benchmark and counterfactual (weights $\upsilon_h$, $\omega_h$ summing to one; weights differ because reweighting changes them). Consistency with the CGE requires

$$\hat{Y}' \;=\; Y'\,\frac{\hat{Y}}{Y}, \tag{5}$$

imposed distribution-neutrally by a single scalar $\lambda$,

$$\lambda \sum_h \omega_h\,\widehat{y}_h \;=\; Y'\,\frac{\hat{Y}}{Y}. \tag{6}$$

Applying $\lambda$ *alone* to the reweighted survey—without the premia update—is the "distribution-neutral growth" treatment of Figure 3. Because $\lambda$ scales all income uniformly, it is also the only channel through which non-labor income (capital, land, remittances, transfers) moves: GIDD does not separately micro-update those components. As Bussolo, De Hoyos, and Medvedev (2010) note, the approximation is accurate for poor households, which hold little capital, while introducing error for wealthier ones.

## A.5 Step 4 — relative-price (food / non-food) adjustment

Real per-capita income deflates monetary income by a household-specific price index reflecting the household's food budget share $\alpha_h$:

$$y^c_h \;=\; \frac{y_h}{\alpha_h P_f + (1-\alpha_h)P_{nf}}, \tag{7}$$

with $P_f, P_{nf}$ the food and non-food price indices from the CGE. The food share is estimated on the survey as

$$\alpha_h = \beta_0 + \beta_1 \ln(y_h) + e_h, \tag{8}$$

and re-evaluated at counterfactual income to give $\hat{\alpha}_h$, which enters eq. (7) in the counterfactual. Equation (8) is the Engel relationship plotted in Figure 1.

## A.6 Multi-sector skill×gender extension

The published derivation above uses two sectors and updates returns analytically. The implementation released with this paper generalizes it in three ways, applied between Steps 1 and 2 above:

1. **Labor classification.** Earners are typed by skill (two levels, education-proxied) crossed with gender, giving four labor types, each observed across the model's full set of production sectors.

2. **Sectoral reallocation.** For each skill–gender type, the CGE's sectoral labor-demand changes define expanding and contracting sectors. Workers are reassigned from contracting to expanding sectors up to the implied quota; the selection among candidates uses an estimated multinomial-logit of sector choice on worker characteristics (with region among the covariates), so reallocation follows observed mobility patterns rather than being random.

3. **Mincer imputation for movers.** A relocated worker's wage in the receiving sector is predicted from a sector-specific Mincer equation (log-earnings on age, age$^2$, education, education$^2$); the residual is rescaled by the ratio of receiving- to origin-sector residual standard deviations, preserving within-sector wage dispersion. The premia reconciliation of eq. (4) is then applied on the skill$\times$gender$\times$sector margin.

Rural/urban location enters only as a covariate in the multinomial-logit of step 2, not as a separate wage cell; a spatially resolved earnings update is left to future work (§7).

## A.7 Replication code

The Stata implementation of the algorithm above is released with this paper in the `gidd/` directory of the project repository. It comprises the driver `gidd.ado` and its components—demographic reweighting (`reweight.ado`), CGE import (`giddcgeread.ado`, `giddgdximport.ado`), sectoral reallocation (`giddsecmove.ado`), Mincer imputation (`giddmincer.ado`), the premia and income scaling of Steps 2–3 (`giddecon3.ado`), and the orchestration script (`giddmaster.ado`). The repository `README` maps each numbered step above to the corresponding routine. Compiled binaries and licensed population data are excluded; the code is provided for transparency of the algorithm, not as a turnkey run.